

0300
PATENT APPLICATION #3

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

Internal Application of

Paul ANUZIS et al.

Application No.: 09/898,008

Filed: July 5, 2001

Docket No.: 110023

For HEALTH MONITORING

CLAIM FOR PRIORITY

Director of the U.S. Patent and Trademark Office
Washington, D.C. 20231

Sir:

The benefit of the filing date of the following prior foreign application filed in the following foreign country is hereby requested for the above-identified patent application and the priority provided in 35 U.S.C. §119 is hereby claimed:

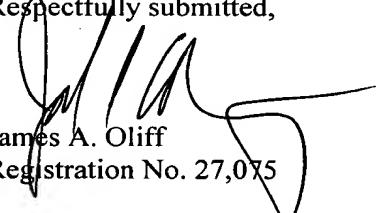
Great Britain Application No. 0016561.3 filed July 5, 2000.

In support of this claim, a certified copy of said original foreign application:

 X is filed herewith.
 was filed on in Parent Application No. filed .
 will be filed at a later date.

It is requested that the file of this application be marked to indicate that the requirements of 35 U.S.C. §119 have been fulfilled and that the Patent and Trademark Office kindly acknowledge receipt of this document.

Respectfully submitted,


James A. Oliff
Registration No. 27,075

Joel S. Armstrong
Registration No. 36,430

JAO:JSA/zmc

Date: August 23, 2001

OLIFF & BERRIDGE, PLC
P.O. Box 19928
Alexandria, Virginia 22320
Telephone: (703) 836-6400

DEPOSIT ACCOUNT USE
AUTHORIZATION
Please grant any extension
necessary for entry;
Charge any fee due to our
Deposit Account No. 15-0461



INVESTOR IN PEOPLE

The Patent Office
Concept House
Cardiff Road
Newport
South Wales
NP10 8QQ



I, the undersigned, being an officer duly authorised in accordance with Section 74(1) and (4) of the Deregulation & Contracting Out Act 1994, to sign and issue certificates on behalf of the Comptroller-General, hereby certify that annexed hereto is a true copy of the documents as originally filed in connection with the patent application identified therein.

In accordance with the Patents (Companies Re-registration) Rules 1982, if a company named in this certificate and any accompanying documents has re-registered under the Companies Act 1980 with the same name as that with which it was registered immediately before re-registration save for the substitution as, or inclusion as, the last part of the name of the words "public limited company" or their equivalents in Welsh, references to the name of the company in this certificate and any accompanying documents shall be treated as references to the name with which it is so re-registered.

In accordance with the rules, the words "public limited company" may be replaced by p.l.c., plc, P.L.C. or PLC.

Re-registration under the Companies Act does not constitute a new legal entity but merely subjects the company to certain additional company law rules.

Signed

Dated 18 July 2001

OLIFF & BERRIDGE, PLC
P.O. BOX 19928
ALEXANDRIA, VA 22320
(703) 836-6400
APPLICANT: Paul ANUZIS et al.
APPLICATION NO.: 09/898,008
FILED: July 5, 2001
FOR: HEALTH MONITORING
ATTORNEY DOCKET NO.: 110023

Request for grant of a patent

(See the notes on the back of this form. You can also get an explanatory leaflet from the Patent Office to help you fill in this form)



The Patent Office

Cardiff Road
Newport
Gwent NP9 1RH

1. Your reference

2. Patent application number

(The Patent Office will fill in this part)

0016561.3

05 JUL 2000

3. Full name, address and postcode of the or of each applicant (underline all surnames)

Kolls-Royce plc
65 BUCKINGHAM GATE
LONDON
SW1E 6AT

Patents ADP number (if you know it)

397000 2

If the applicant is a corporate body, give the country/state of its incorporation

ENGLAND

4. Title of the invention

HEALTH MONITORING

5. Name of your agent (if you have one)

MEWBURN ELLIS

"Address for service" in the United Kingdom to which all correspondence should be sent (including the postcode)

YORK HOUSE
23 KINGSWAY
LONDON
WC2B 6HP

Patents ADP number (if you know it)

109006 ✓

6. If you are declaring priority from one or more earlier patent applications, give the country and the date of filing of the or of each of these earlier applications and (if you know it) the or each application number

Country

Priority application number
(if you know it)

Date of filing
(day / month / year)

7. If this application is divided or otherwise derived from an earlier UK application, give the number and the filing date of the earlier application

Number of earlier application

Date of filing
(day / month / year)

8. Is a statement of inventorship and of right to grant of a patent required in support of this request? (Answer 'Yes' if:

YES

a) any applicant named in part 3 is not an inventor, or
b) there is an inventor who is not named as an applicant,
or

c) any named applicant is a corporate body.

See note (d))

9. Enter the number of sheets for any of the following items you are filing with this form. Do not count copies of the same document

Continuation sheets of this form 0

Description 18

Claim(s) 0

Abstract 0

Drawing(s) 2

10. If you are also filing any of the following, state

Priority documents 0

Translations of priority documents 0

Statement of inventorship and right to grant of a patent (Patents Form 7/77) 0

Request for preliminary examination and search (Patents Form 9/77) 0

Request for substantive examination (Patents Form 10/77) 0

Any other documents (Please specify) 0

11. I/We request the grant of a patent on the basis of this application.

Signature

Stephen Carter

Date

5 July 2000

12. Name and daytime telephone number of person to contact in the United Kingdom STEPHEN CARTER 020 7240 4405

Warning

After an application for a patent has been filed, the Comptroller of the Patent Office will consider whether publication or communication of the invention should be prohibited or restricted under Section 22 of the Patents Act 1977. You will be informed if it is necessary to prohibit or restrict your invention in this way. Furthermore, if you live in the United Kingdom, Section 23 of the Patents Act 1977 stops you from applying for a patent abroad without first getting written permission from the Patent Office unless an application has been filed at least 6 weeks beforehand in the United Kingdom for a patent for the same invention and either no direction prohibiting publication or communication has been given, or any such direction has been revoked.

Notes

- If you need help to fill in this form or you have any questions, please contact the Patent Office on 0645 500505.
- Write your answers in capital letters using black ink or you may type them.
- If there is not enough space for all the relevant details on any part of this form, please continue on a separate sheet of paper and write "see continuation sheet" in the relevant part(s). Any continuation sheet should be attached to this form.
- If you have answered 'Yes' Patents Form 7/77 will need to be filed.
- Once you have filled in the form you must remember to sign and date it.
- For details of the fee and ways to pay please contact the Patent Office.

HEALTH MONITORING

Field of the Invention

This invention relates to methods and data processing systems for monitoring the health of a system.

5 The methods and data processing systems of the invention are particularly, although not necessarily exclusively, suitable for monitoring the health of power plant, including for example gas turbine, spark ignition and compression ignition internal combustion engines.

Background

10 The health of a system can be considered a measure of the condition of a system against expected norms. A healthy system is one whose condition closely matches expectations, whereas an unhealthy system is one
15 whose condition differs from what would be expected, indicating for example deterioration of, or a possible problem with the system. The ability to monitor the health of a system can therefore allow such deterioration and/or problems to be detected and, if necessary,
20 addressed at an early stage.

In order to determine the condition, and consequently health, of a system, it is normal to monitor and analyse a series of measurable indicators which themselves reflect aspects of the condition of the
25 system. For instance, taking the example of a gas turbine, one might monitor performance parameters such as

turbine and compressor operating temperatures and pressures and spool speeds. To obtain a fuller overall picture of the engine's condition, these performance parameters can be supplemented with further condition indicators including, for example, vibration measurements and a measure of the amount of debris in the circulating oil.

Particularly with complex mechanical / thermodynamic systems such as gas turbines and other internal combustion power plant, the number of indicators that must be monitored to obtain a useful overall picture of the system's condition can be high. This in turn means that the task of analysing the complete series of indicators to determine the health of the engine is a complex one, typically requiring a skilled expert to analyse the data off-line.

Taking again the example of a gas turbine, it is known for example to collect performance and vibration data from the engine over time to be analysed off-line by one or more experts. Typically the performance data will be compared with simulated data for the same engine and, based on this comparison, an expert will form a view as to the health of the engine. Additionally, the vibration data will be reviewed, although often only superficially. If a problem is detected, the vibration data may then be analysed in more detail, often by another expert, to look for any unusual patterns which might indicate a loss of health.

Summary of the Invention

It is a general aim of the present invention to provide methods and data processing systems that facilitate the acquisition and analysis of condition indicators in a manner such that the overall health of a system can be more readily assessed.

Accordingly, in a first aspect, the invention provides a method for monitoring the health of a system, comprising:

constructing a system condition signature from a plurality of measured condition indicators acquired from the system;

comparing the system condition signature with a normal signature, corresponding to the signature for a healthy system; and

registering an event if the system condition signature differs from the normal signature by more than a predetermined threshold.

By merging or fusing the condition indicators into a single signature in this manner, and providing a normal signature with which this fused data can be compared, the task of assessing the health of a system is greatly simplified. In particular, since the detection of an event amounts to an indication of an unhealthy system (i.e. a system condition that differs from what would normally be expected), the monitoring of the health can be largely automated, removing, or at least minimising, the requirement for expert input during the

monitoring process. This in turn means that it becomes feasible to continuously monitor the health of a system, and to provide useful information about the health of the system in real time during operation.

5 Preferably the condition indicators that are combined to form the system condition signature include performance parameters, which in the case of a mechanical system may be speeds, pressures and temperatures for example. Other useful parameters may include what might
10 be conventionally thought of as control or status parameters. For convenience, such parameters will be referred to using the single label of "performance parameters" in the following text.

 Additionally, to obtain a fuller picture of a
15 mechanical system's health, it is particularly preferred that the signature includes one or more condition indicators related to the vibration of the system.

 Put, more generally, the condition indicators from which the system condition signature is constructed
20 may be derived from two or more disparate sources of data. This illustrates a particular strength of this approach in that a great variety of different forms of condition indicator data can be encompassed in the system condition signature, providing a more comprehensive
25 measure of the system's health than has previously been possible without multiple analyses.

 The normal signature for the system can be derived from a predefined model of the system that is

being monitored. This model can itself be developed off-line and then fixed for the duration of the operation of the health monitoring method. More preferably, however, the model is designed to be refined as the method
5 proceeds in order that it might be better tuned to a specific system.

Whichever approach is adopted, it is particularly preferred that the model is developed using a data-driven, or at least partially data-driven approach. That
10 is to say the model learns from training data comprising series of the condition indicators which have been labelled as normal (i.e. healthy) or abnormal (i.e. unhealthy) as the case may be. In fact, it is often the case that normal data is far more readily available than
15 abnormal data and therefore the training data may only include examples of normal data. This still results in an effective model, because subsequent events can then be identified as departures from the trained model of normality.

20 The invention also provides a data processing system for monitoring the health of a system, comprising:

data acquisition means for acquiring a plurality of measured condition indicators from the system;

processor means for constructing a system
25 condition signature from said plurality of measured condition indicators;

comparator means for comparing the system condition signature with a predefined normal signature,

corresponding to the signature for a healthy system; and
means for registering an event if the comparator
indicates that the system condition signature differs
from the normal signature by more than a predetermined
5 threshold.

Where the system condition signature is comprised
of data from disparate sources, for instance performance
and vibration data, a problem occurs in that the data may
well not be synchronised in time. If this asynchronous
10 data is combined to form the signature, a distorted
picture of the system's health may well result. For
similar reasons, training data used to develop a model of
normal system behaviour should also be synchronised if
distortions in the model are to be avoided.

15 To address this problem, the invention proposes
to associate time stamps (based on a common clock) with
the acquired data and to synchronise the data on the
basis of these time stamps.

Accordingly, in a second aspect, the invention
20 provides a method of synchronising two or more data
streams, each data stream comprising a series of
sequentially acquired data elements, the method
comprising:

associating a time stamp with each data element
25 of each stream, the time stamp identifying the time of
acquisition of the data element on the basis of a clock
common to all data streams;

selecting a first data element from a first

stream and inspecting its time stamp;

conducting a search of the data elements of the
or each other stream to identify the data element in the
or each other element having an associated time stamp
5 closest to that of the selected element of the first
stream; and

marking said identified data element of the or
each other stream and said selected element of the first
stream as being synchronised with one another.

10 The process can be repeated until the data
elements in the first stream have been exhausted. In any
subsequent processing of the data that is reliant on
using synchronised data streams, only those data elements
marked as being synchronised with one another are used.

15 In the case where the acquisition rates of the
data streams differ from one another, it is preferred
that the first stream, with which the other streams are
synchronised, is chosen to be the stream having the
lowest acquisition rate.

20 The invention further provides a data processing
system for synchronising two or more data streams, each
data stream comprising a series of sequentially acquired
data elements, comprising:

means for associating a time stamp with each data
25 element of each stream, the time stamp identifying the
time of acquisition of the data element on the basis of a
clock common to all data streams;

means for selecting a first data element from a

first stream and inspecting its time stamp;

means for conducting a search of the data elements of the or each other stream to identify the data element in the or each other element having an associated time stamp closest to that of the selected element of the first stream; and

means for marking said identified data element of the or each other stream and said selected element of the first stream as being synchronised with one another.

10 **Brief Description of the Drawings**

The various aspects of the invention will be further described by way of example with reference to the accompanying drawings, in which:

15 Fig. 1 illustrates the basic strategy underlying the first aspect of the invention;

Fig. 2 schematically illustrates an exemplary data structure that can be adopted for operation of the second aspect of the invention;

20 Fig. 3 shows the learning curve for a system model employing the preferred hybrid approach; and

Fig. 4 shows a comparison of observations and modelled estimates for a shaft speed measurement, illustrating evolution of the preferred model.

Description of the Embodiments

25 The embodiment described below is an example of a data processing system employing both aspects of the

invention discussed above. More specifically, it is a system for synchronous acquisition, analysis and display of performance parameters and vibration data from a power plant (e.g. a gas turbine), for monitoring the health of the plant.

In accordance with a preferred form of the second aspect, the performance and vibration data streams are synchronised in real time and, in accordance with a preferred aspect of the first aspect of the invention, these data are combined or fused to construct a signature for the system that can be compared to a signature derived from a model representing a healthy power plant, in order to provide anomaly/event detection and hence fault diagnosis.

The following discussion focusses on an application of the system to monitoring the health of a gas turbine aero-engine, but it will be appreciated that the methods can be adapted to other power plant, including for example ground-based and marine gas turbines, and spark ignition and compression ignition internal combustion engines, as well as other mechanical, thermodynamic, fluid, electrical or electronic systems for example.

The system acquires performance parameters from the gas turbine digitally via an ethernet link at a rate between 20 and 40Hz. Typical performance parameters are measurements of pressure, temperature, thrust, altitude or Mach number. Vibration data is acquired from analogue

vibration transducers which are sampled at user-selectable sampling rates (from 625Hz to 80kHz) via an analogue-to-digital converter. The amplitude spectrum of the vibration data is generated using the Fast Fourier Transform every 0.2 sec.

The performance and vibration data streams are asynchronous and stored in separate files together with the corresponding timestamps. During review, as data is loaded into memory, synchronisation is performed between the performance and spectrum data on a line by line basis. Markers 10,12 (see Fig. 2) are kept which record the last synchronised line in the vibration and performance data ring buffers 14,16. When new data is available in memory, the timestamp for the next vibration spectrum line is examined. The synchronisation algorithm starts from the last previously synchronised location in the performance data and searches forwards or backwards based on the timestamps of the performance data (accurate to 0.05 sec) until the closest matching timestamp in the performance data ring buffer 16 is identified. This location in the performance data is recorded as being synchronised with the line in the vibration ring buffer 14. The algorithm then proceeds to the next line in the vibration ring buffer 14 (0.2 sec later) and so on until there is no more data available to synchronise.

Considering the synchronisation algorithm in a little more detail, it can be seen from Fig. 2 that the algorithm maintains a synchronisation table 18 that gives

the index of the performance data entry that matches each vibration data line. The algorithm uses variables to mark the latest synchronised data in each buffer. The operation of the algorithm can be summarised by the following 'pseudo code':

1. Initialise the latest synchronised markers to the start of the vibration and performance data.
2. Loop while there is more data in both ring buffers.
 - (a) Starting from the latest synchronised data item in each ring buffer, examine the time stamp, t , on the next entry in the vibration ring buffer.
 - (b) Search forward in the performance ring buffer until a time stamp greater than t is found. Select between this entry in the performance ring buffer and the previous entry for one which is closest to t and record the match in the synchronisation table.

Once synchronised, the analysis of this performance and vibration data relies on constructing models of normal jet engine behaviour and then detecting an event or an abnormality with respect to these models.

5 The overall strategy is shown in Fig. 1.

In change detection, the evolution of a sequence from a signal is examined over time. Related to this is the use of the innovation in Kalman filtering systems (see M. Gelb *Applied Optimal Estimation*. MIT Press 1974
10 for background on this). By contrast, novelty detection is defined here as the examination of a system in the light of information learned, in a prior training phase, from multiple examples of similar systems. All of these processes can be used to detect events or abnormalities.

15 As noted above, this example is concerned with condition monitoring in gas turbine engines. Traditional aircraft engine monitoring systems are based on two distinct processes: the use of vibration signatures to indicate engine state, and a separate procedure, gas-path
20 analysis, which is employed for determination of state from performance parameters. In the approach described now, however, performance-related parameters such as pressure and temperature are fused with vibration data. Furthermore, it is proposed to employ a hybrid of
25 knowledge-based and data-driven models to model a normal engine.

The aim is to take advantage of disparate sources of data to form a more comprehensive picture of engine

state during normal operation. This in turn should allow a wider range of deviations to be identified. In addition, fusing two methods of data-analysis allows the accuracy of prior expert knowledge to be combined with the robustness of data-driven approaches. Thus, although models of the engine system are used, these are not fixed. Instead, they are data-driven models, which evolve with acquired data. This offers the important advantage of robustness.

Two alternative methods of data analysis methods are described below. They are distinguished by the amount of prior knowledge required to set up the system. In both cases, the role of the expert need only be retained in classifying training data as novel or normal.

Novelty Detection

The first method is referred to as Novelty Detection (box A in Fig. 1). This approach relies on a prior learnt model of normality. For example, normal vibration tracked order shapes are learnt using a simple clustering model for the normal data. The novelty of eg. the vibration signature for an engine under test is assessed by comparing the closeness of its tracked order signature with the prototypical patterns in the clustering model of normality. This can be done, for example, by computing the shortest normalised Euclidean distance between the vector encoding the tracked order shaped to any of the (prototypical patterns) cluster

centres in the model of normality (see Nairac et al, "A System for the Analysis of Jet Engine Vibration Data", *Integrated Computer-Aided Engineering*, 6(1):53-65, 1999).

If this distance is beyond a previously set threshold, the vibration signature as represented by that tracked order is deemed to be outside the bounds of normality.

In addition to the vibration tracked orders, the model of normality for the vibration spectra includes the following: sidebands, multiple harmonics, fractional harmonics and broadband power.

Innovation Detection

The other data analysis method used for event and abnormality detection is the use of process and observation models (in a Kalman filter, for example) for innovation (box B in Fig. 1). The process model has a state vector associated with it (see below). The observation vector in this case includes both performance parameters and vibration information so that two types of data are fused within the Kalman filter model. The fusion of the data is performed in real-time with a new output being generated by the system several times a second.

An important aspect of the use of this model in the system is the use of learning. In a first, off-line, phase of learning, a generic model of the engine is learnt. The learning is data-driven using an algorithm such as Expectation-Maximisation in order to maximise the

likelihood of the model given the data. Once such a generic model has been learnt off-line for a particular type of engine, learning can then be applied on-line in order to tune the model to an individual engine

5 immediately after its pass-off test and after each maintenance procedure. Engine deterioration can also be learnt on-line. The model can be tuned to different flight conditions, such as cruising or landing, in order to identify innovation with even more sensitivity and
10 specificity.

A further aspect of the use of these models in the system is the integration of the data-driven learnt models with existing performance models which rely on the laws of thermodynamics and computational fluid dynamics
15 (knowledge models). These models can therefore be described as hybrid models because they are based on the integration of data-driven and knowledge-based models.

Looking in more detail at the preferred hybrid approach, it is known to apply Expectation Maximisation
20 to parameter estimation in linear dynamical systems and to a limited degree to non-linear systems (see Ghahramani and Hinton. *Parameter estimation for linear dynamical systems*. Technical Report CRG-TR-96-2, University of Toronto, 1996; Roweis and Ghahramani, "A Unifying Review
25 of Linear Gaussian Models", *Neural Computation*, 11, 305-345, 1999; Ghahramani and Roweis, "Learning in nonlinear dynamical systems using an EM algorithm" in Kearns et al (editors) *Advances in Neural Information Processing*

Systems, Volume 11, MIT Press, 1999.). This approach can be used for hybrid model development in the present case.

The EM learning algorithm is applied to a Kalman filter model. In the linear case, this is a system with
5 a measurement process of the form

$$y(i) = Cx(i) + v(i) \quad (1)$$

where $y(i)$ is a set of observations of hidden state $x(i)$ and measurement noise $v(i)$ is zero-mean and normally distributed with covariance matrix R . The state equation
10 is

$$x(i+1) = Ax(i) + w(i) \quad (2)$$

with $w(i)$ zero-mean and normally distributed with covariance matrix Q .

To illustrate this, consider a simple example
15 where observations are made of the speeds of the three shafts of a test engine during cruise. The observed data y is simply the state x corrupted by noise, so

$$y(i) = x(i) + v(i) \quad (3)$$

The observations are used during the learning process, to
20 generate a dynamical system model in which A , C , Q and R are learned from the data. The process and observation models and process and observation noise covariances are initialised randomly.

Fig. 3 illustrates the learning (log likelihood)
25 plot for the system. Fig. 4 shows the evolution of estimates of shaft 1 speeds during the learning process using an on-line version of the EM algorithm [4]. The learning stage lasts for the first 25 iterations. From

iteration 25 onwards, the system evolves according to the learned matrices (which are then kept fixed).

To combine the vibration information with the performance parameters in this hybrid approach, it is possible to augment the performance state vector with vibration parameters. For example, the vibration information can be represented by tracked order vectors - the elements of which correspond to the signal amplitude in a narrow range centred on the main vibration frequencies for each shaft.

Once trained, the systems can be used to detect events or abnormalities, that is to say divergences from the learnt model of normality. The events of particular interest are those that are unexpected, possibly indicating a problem with the engine for example. However, particularly where the models have been learned only for "steady state" parts of the flight envelope (e.g. acceleration, cruise and deceleration), transients during operation of the engine will also be flagged up as events, although they are expected. For example, where a bleed valve is opened or closed, the operating condition of the engine will exhibit significant differences from a learnt model of steady state normality.

Measures are preferably employed to avoid distortion of the model by these transient events. For instance, since the opening of a bleed valve is an event that occurs at a defined point in time, the data collected from the engine at that time and slightly

either side of it (e.g. for 2 seconds before and after) can be eliminated from the data analysed by the health monitoring system.

5 The system also includes a display which is driven to allow information to be displayed either during acquisition or for review once an acquisition cycle has been completed. It preferably includes the following features:

- 10 • Ability to display a combination of any two of vibration spectra, tracked orders, broadband power, performance parameters synchronised in time.
- Ability to extract and plot vibration spectra against engine speed.
- 15 • Ability to interrogate and print any of vibration spectra, tracked orders, broadband power and performance parameters.
- Automatic detection and display of features from vibration spectra (sidebands, harmonics, etc.)

1/2

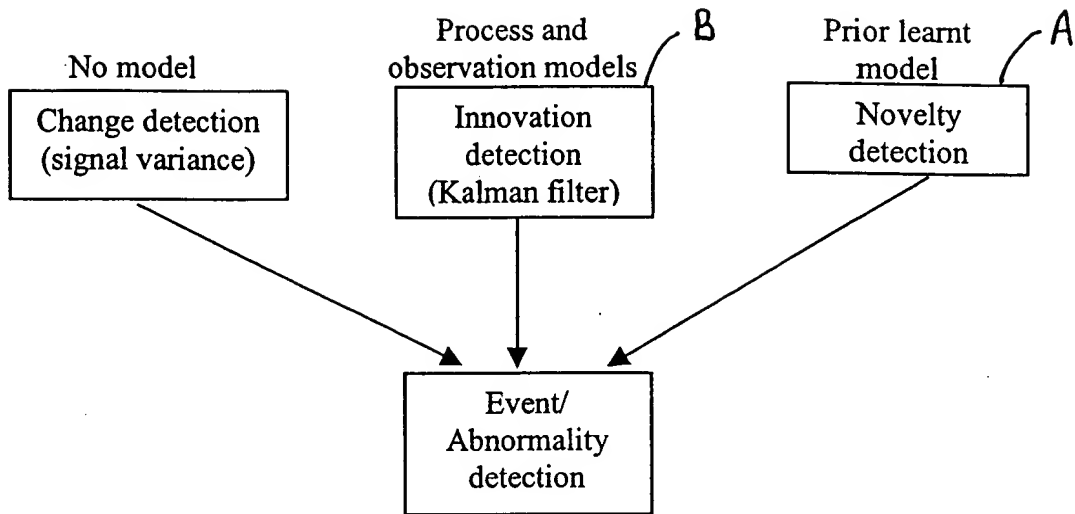


FIG 1.

2/2





FIG 3.




FIG. 4.